

# The Role of Familiarity, Priming and Perception in Similarity Judgments

Laura M. Hiatt and J. Gregory Trafton

U.S. Naval Research Laboratory

Washington, DC 20375 USA

{laura.hiatt, greg.trafton}@nrl.navy.mil

## Abstract

We present a novel way of accounting for similarity judgments. Our approach posits that similarity ratings stem from three main sources: familiarity, priming, and inherent perceptual similarity. We present a process model of our approach in the cognitive architecture ACT-R, and match our model's predictions to data collected from a human subject experiment which involved simple perceptual stimuli. Familiarity accounts for rising ratings over time; priming accounts for asymmetric effects that arise when the stimuli are shown with different frequencies. Pure perceptual similarity also predicts trends in the results. Overall, our model matched the data with  $R^2$  of 0.99.

## Introduction

Similarity is a critical and pervasive part of human cognition (Medin, Goldstone, & Gentner, 1993). Similarity measures are integral to object categorization and classification (Nosofsky, 1992). Similarity is also pervasive in problem solving (Novick, 1990), decision-making (Medin, Goldstone, & Markman, 1995), and memory (Roediger, 1990). As with many aspects of human cognition, however, the mechanisms that determine similarity are not yet fully understood. Various theories abound, with none yet able to capture enough different types of situations to be called the winner (Rorissa, 2005).

One interesting result in this field is asymmetries that have been shown to arise when making similarity judgments, even of very simple perceptual stimuli (Tversky, 1977; Rosch, 1975). Rosch (1975) argued that such similarity is based on mapping stimuli onto one another and, intuitively, non-prototypical stimuli map more easily onto prototypical stimuli than vice versa. Tversky (1977) argued that is due to weighted feature matching, where the salience of features in the current context determines their weight; others agree with this thought in general (Medin et al., 1993; Glucksberg & Keysar, 1990).

These two explanations, however, assume that either there is a clear prototype inherently present in the experiment (such as the more perceptually complex stimulus), or that stimuli have various features which have a clear inherent order of cognitive preference and saliency (such as symmetry). They do not, however, provide any concrete explanations for why complexity or symmetry may lead to prototypicality or saliency.

Polk, Behensky, Gonzalez, and Smith (2002) shed light on the situation by presenting an experiment that avoids the question entirely by using perceptual stimuli where the only feature was color (so there are no features to comparatively weigh), and where the color hues are fairly similar (so there is

no clear prototype). The experiment showed a striking asymmetry in similarity judgments between the different colors when they were presented with different frequencies during an irrelevant training task: colors which had been trained on *less* often were considered more similar to colors which had been trained on *more* often than the other way around. To account for these low-level results, Polk et al. (2002) implemented a neural network which simulated the asymmetry by measuring the ease with which the network switches between different activation patterns; those that are more stable (e.g., high-frequency patterns) were easier to assimilate to.

In our approach, we match the human subject data from Polk et al. (2002) while attempting to address three additional points. First, there was a second significant effect, that the ratings in general increased over time, that the above models do not address. Second, we wanted our approach to provide explicit cognitive processes for similarity ratings. Third, we believe that inherent perceptual similarity also plays a part in these types of similarity judgments (e.g., purple is inherently more similar to blue than to orange) (Smith & Heise, 1992).

To this end, we begin our approach with the cognitive architecture ACT-R (Adaptive Characterization of Thought – Rational) (Anderson, 2007). Using ACT-R, we account for similarity judgments by considering three values provided *a priori* by the architecture. The first, familiarity, is represented as a base-level activation value of a concept, which represents its frequency and recency of use. The second, priming, is based on spreading activation, which disperses activation between different, associated concepts in declarative memory (Anderson, 1983; Harrison & Trafton, 2010). In addition, we utilize an extension to ACT-R which provides it with a calculation for measuring color similarity (Breslow, Ratwani, & Trafton, 2009; Breslow, Trafton, & Ratwani, 2009).

Our model starts without any pre-existing declarative knowledge or network structures; all knowledge is created during the experiment. Over time, our cognitive model builds a network of concepts (e.g., color blocks) by learning associations between them in the form of subsymbolic connections between their representations in declarative knowledge. During each similarity judgment, the model combines its measure of perceptual similarity with base-level and spreading activation to determine its response. On the first trial of an experiment, there is no spreading activation since there is no declarative knowledge and so the judgment is based purely on base-level activation and perceptual similarity; however, over time the model builds up declarative memories that may contribute to spreading activation in later trials. This explains the two main effects found in Polk et al.'s experiment. Dur-

Report Documentation Page			Form Approved OMB No. 0704-0188	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE <b>AUG 2013</b>	2. REPORT TYPE	3. DATES COVERED <b>00-00-2013 to 00-00-2013</b>		
4. TITLE AND SUBTITLE <b>The Role of Familiarity, Priming and Perception in Similarity Judgments</b>			5a. CONTRACT NUMBER	5b. GRANT NUMBER
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	5e. TASK NUMBER
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Naval Research Laboratory, 4555 Overlook Ave., SW, Washington, DC, 20375</b>			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>				
13. SUPPLEMENTARY NOTES <b>in Proceedings of the 35th Annual Meeting of the Cognitive Science Society, Berlin, Germany, 31 Jul ? 3 Aug 2013.</b>				
14. ABSTRACT <b>We present a novel way of accounting for similarity judgments. Our approach posits that similarity ratings stem from three main sources: familiarity, priming, and inherent perceptual similarity. We present a process model of our approach in the cognitive architecture ACT-R, and match our model's predictions to data collected from a human subject experiment which involved simple perceptual stimuli. Familiarity accounts for rising ratings over time; priming accounts for asymmetric effects that arise when the stimuli are shown with different frequencies. Pure perceptual similarity also predicts trends in the results. Overall, our model matched the data with R2 of 0.99.</b>				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES <b>6</b>
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>		

ing judgments made on stimuli which have been previously viewed with different frequencies, less activation is spread from high- to low-frequency color patches than from low- to high-frequency color patches because that is the direction which priming favors. In addition, base-level activation is higher at the end of the experiment than in the beginning due to increased familiarity with the colors, leading to increased ratings over time.

The primary contribution of this work is a general account for similarity which provides implemented, explicit, process-level mechanisms for calculating similarity values. Other work has used activation for similarity in more abstract terms: the neural network written by Polk et al. (2002) relies on activation patterns; and other accounts also exist (Kozima & Furugori, 1993; Ulhaque & Bahn, 1992). Tversky (1977)'s discussion of salience, and Rosch (1975)'s on prototypicality can also be seen as broadly touching upon activation in similarity. Our work solidifies these accounts in a cognitive setting by positing that the abstract notion of salience (or prototypicality) translates to familiarity and priming in a cognitive model. In addition, our work is distinguished because we also account for inherent perceptual similarity. We show that our model is an excellent fit for empirical human subject data on similarity judgments on simple perceptual stimuli.

## Experiment

There were three phases to the experiment: a pre-test phase, a training phase, and a post-test phase (Polk et al., 2002). In the pre-test phase, participants viewed two patches of different colors that were the same size and were asked to rate their similarity on a scale of 0 to 9 (0 as highly dissimilar, 9 as highly similar). Five different hues of green and five different hues of blue were used, designated as *blue1*...*blue5*, and *green1*...*green5*. Greens and blues were never compared to each other; only hues of the same color were shown concurrently. During a trial, the stimuli were presented as part of a text question to emphasize the directionality of the judgment: "How similar is (color patch 1) to (color patch 2)?" The color blocks were labeled "Blue1" and "Blue2," or "Green1" and "Green2" as appropriate, with the label displaying below the color patches; note that this is irrespective of whether the color itself was *blue1*, *blue5*, etc. The blocks were 140X140 pixels. Once a user entered a rating, the screen was cleared for 500ms before the next comparison appeared. Each pair of colors was presented twice in each direction for a total of four times each. The order in which the pairs were presented was randomized, except that the same hue was not present in consecutive trials. The sentence was centered both horizontally and vertically.

In the training phase, participants saw two patches of the same color but different sizes (125X125, 131X131, 138X138 and 144X144 pixels, appearing with equal probability) and were asked to specify which was larger. The key part of the experiment is that, during this phase, two of the five greens and two of the five blues were presented ten times more frequently

as others, 110 times instead of 11 (half the participants saw *blue1*, *blue2*, *green1* and *green2* with a higher frequency, called the "1-2 group", the other saw *blue4*, *blue5*, *green4* and *green5* presented more often, the "4-5 group").

The third phase was a second testing phase that was an exact repeat of the first phase. Forty-five participants took part in the experiment, with ten being excluded due to inaccuracy on their size judgments or self-reporting of a lack of concentration. For more details, see (Polk et al., 2002).

The experiment produced two interesting results. First, there was no significant difference between "forward" (less frequent color on the left, more frequent color on the right) and "backward" (more frequent color on the left, less frequent on the right) comparisons in the pre-test phase; however, these comparisons showed a striking asymmetry effect during the post-test phase. Specifically, forward comparisons were ranked as significantly more similar than backward comparisons during this testing phase. A second find was that similarity ratings were significantly higher in the post-test than in the pre-test. No other effects were reported as significant by the authors.

## ACT-R

At the core of our approach is the cognitive architecture ACT-R (Anderson, 2007). At a high level, ACT-R is a hybrid symbolic/subsymbolic production-based system. In other words, given declarative knowledge (fact-based memories, or "chunks") and procedural knowledge (rule-based memories, or "productions"), as well as input from the world (e.g., visual), it decides what productions to fire next; these productions can either change its internal state (e.g., by creating new knowledge) or its physical one (e.g., by pressing a key on a keyboard). Knowledge has both a symbolic component, such as who was where at what time, and subsymbolic one, such as how relevant a fact is to the current situation.

ACT-R is made of up several major components. First, it has several limited-capacity buffers. Each buffer is backed by one (or more) theoretically motivated modules (e.g., declarative, visual, aural, etc.); in addition, there is the procedural module, which does not have an associated buffer. Each module represents a specific cognitive faculty and has been shown to have anatomical correspondences in the brain (Anderson, Albert, & Fincham, 2005; Anderson, 2007).

Chunks consist of a set of slots, whose values determine the concept that the chunk represents. At any point in time, there may be at most one chunk in any individual buffer; a module's job is to decide when to put chunks into its corresponding buffer. Then, a central pattern matcher uses the contents of the buffers, if any, to match specific productions which, when fired, can modify the current buffer contents.

The relevant modules of ACT-R to this paper are the declarative, intentional, imaginal, visual and motor modules, which are associated with the retrieval, goal, imaginal, visual and visual-location, and motor buffers, respectively. The declarative module manages the creation and storage of the model's

factual memory; in addition, when requested, chunks can be accessed via the retrieval buffer. It has been shown to be an astonishingly good predictor of human declarative memory (Anderson, Bothell, Lebiere, & Matessa, 1998; Anderson, 1983; Schneider & Anderson, 2011). The intentional and imaginal modules provide support for task-oriented cognition. The goal buffer (associated with the intentional module) typically contains chunks that identify and placekeep the model's current goal; the imaginal module usually contains intermediate problem state representations. Finally, the visual and motor modules interface ACT-R with the world, allowing ACT-R to see objects on a computer screen (including their height and color) and press buttons on a keyboard.

To quantify the perceptual difference between the color blocks, we consider a measure of color similarity proposed by Breslow, Ratwani, and Trafton (2009). They introduced a component to ACT-R which supports high-level color processing that can detect both color similarity and brightness difference between colors. It is based on the CIEDE2000 algorithm (CIE, 2001), and has been shown to match well with human subject data. Given two color values, this color similarity component returns a numeric measure of how similar they are perceived to be by the cognitive model.

### Subsymbolic Information

A key aspect of declarative memory in ACT-R is priming, or the subsymbolic activation of chunks. Activation consists of three primary components: base-level activation, spreading activation, and noise. Base-level activation is a measure of familiarity that is learned over time and is a function of the frequency and recency of references to the chunk, where a single reference is defined as (for purposes of this paper) being added to and then removed from a buffer. It is designed to represent the activation of a chunk over longer periods of time. Spreading activation, on the other hand, is temporary and based on the current context, allowing chunks that are the focus of attention to prime related memories for short periods of time. Noise is a random component that models the noise of the human brain; since its presence would not affect our results, we ignore noise in the rest of this paper.

Spreading activation is spread along subsymbolic associative links between chunks. Links are created from a chunk  $j$  to a chunk  $i$  when: (1) chunk  $i$  contains chunk  $j$ , or has chunk  $j$  as one of its slot values; or (2) chunks  $i$  and  $j$  are both matched by the same production (called *co-occurrence*) (Anderson, 1983; Harrison & Trafton, 2010). There are other ways of creating links, as well, but we do not utilize them in this model and so forgo their discussion. Once established, links have an associated strength value which affects how much activation is passed along the link from chunk  $j$  to chunk  $i$ . Link strengths, intuitively, reflect the probability that chunk  $i$  will be needed when chunk  $j$  is being referred to by a production. They are a function of how many times chunks  $j$  and  $i$  have been referred to by a production at the same time, vs. how many times chunk  $j$  was referred to by a production without chunk  $i$ . Note, then, that while links

stemming from co-occurrence are always present in both directions (i.e., chunk  $j$  activates chunk  $i$  and vice versa), the links may be of different strengths if the chunks have not always been referred to by productions at the same time, or with the same frequency.

Spreading activation sources from the goal buffer. When a chunk  $i$  is in the goal buffer, the buffer's source activation is divided equally among all chunks  $j$  which have an outgoing link to chunk  $i$  (such as a slot value of chunk  $i$ , or a chunk that has co-occurred with chunk  $i$  in the past). The  $j$  chunks then use their source activation as the basis of spreading activation along all of their outgoing links. Note that in ACT-R, this is a one-step process; the chunks that receive spreading activation from the  $j$  chunks do not, in turn, spread activation along their outgoing links (Anderson, 1983).

### Model and Fit

The model itself is fairly simple. It starts out with no declarative knowledge, but with the productions necessary to complete the experimental task. For each trial during the two testing phases, the model starts by looking at the color block on the left. Once the model has the block chunk in its visual buffer, the model requests a retrieval of the chunk of the color associated with it, and looks for the object on the right. Once the left color chunk has been retrieved, the model places it in the imaginal buffer and removes it from the retrieval buffer. Then, when it sees the object on the right, it can retrieve the color of the second color block. When that chunk has been retrieved, the model has each of the color chunks in a buffer and it can then proceed to making the similarity judgment.

The model draws similarity from three sources: how perceptually similar the two colors are to each other, how familiar the right color is, and to what degree the right color is primed. Recall that mathematically, perceptual similarity is calculated from the RGB values of two colors; familiarity is represented as base-level activation, which numerically represents the recency and frequency of use of a color; and, finally, priming is measured via the amount of spreading activation that a color receives in the current context. Therefore, the model calculates the perceptual similarity of the two colors, and looks at the total activation (both base-level and spreading) that the right color has at the time of the judgement. Note that while the right color receives spreading activation from sources other than the left color, the amount of "other" spreading activation is constant across trials, making the spreading activation from the left color the cause of asymmetry effects. After the model has these two numbers available, it presses the button corresponding to its rating, the trial is finished and the model waits for the next one to begin.

During a training trial, the model first looks at the color block on the left. It then stores the block's color as part of the goal representation in case it is later needed (such as if the subsequent visual search fails), and looks for another block of the same color. Once it sees the second color block, it retrieves the first one. Then, while thinking about the first

block and looking at the second one, it compares their heights and responds accordingly.

Subsymbolically, during each testing trial, co-occurrence links are created (or strengthened) between: each of the block chunks and the current goal chunk; each of the color chunks and the current goal chunk; and the two color chunks. Block chunks also have an incoming, containment link from their associated color chunk. During a training trial, co-occurrence links are created (or strengthened) between: the two block chunks; and each of the block chunks and the current goal chunk. The goal chunk and the block chunks also each have an incoming link from the color chunk because they contain it as a slot value. Figure 1 shows this in diagrammatic form, showing co-occurrence links as bi-directional for simplicity.

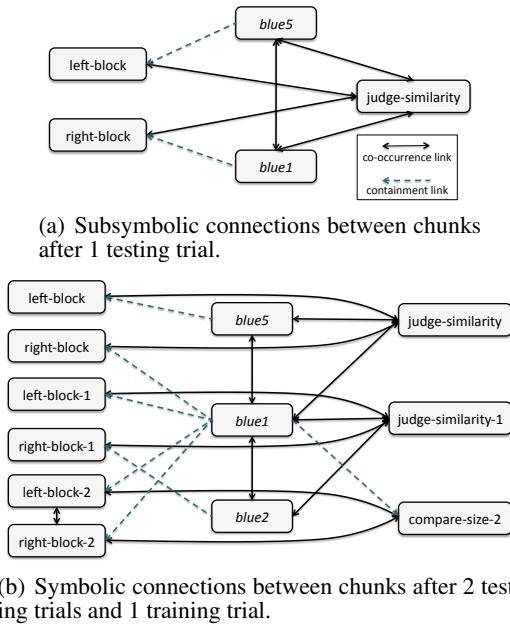


Figure 1: Subsymbolic connections between chunks at various phases of model execution. Here, the model performed two pre-test trials (with colors *blue5/blue1*, and *blue1/blue2*, respectively), and one training trial (where *blue1* is the color). In order to maintain clarity, this diagram is slightly simplified from the model’s actual subsymbolic network (e.g., it does not contain containment links for the visual location slots of the left and right blocks, which have no bearing in the spreading activation process here).

In terms of parameters, the associative learning rate, which affects the rate at which links are strengthened, was set to 6.5, which represents a fairly brisk rate of learning. There is no standard value for this parameter. The base level learning decay parameter was set to 0.4 instead of its default of 0.5. All other parameters were set to their default values.

## Model Predictions

First, the model predicts that later comparisons will, overall, be more similar than earlier comparisons. Before the exper-

iment begins, the colors are not familiar to the model, and so do not have very high base-level activations. During the pre-test, those values increase as the color chunks are referenced many times. Throughout the training phase the chunks’ base-level activations decay, since the color chunks are not referenced in those productions. Base-level activation values then increase again during the post-test, leading to higher familiarity with the colors in the post-test than in the pre-test. Since all colors are shown equally during the pre-test and post-test phases, base-level activation does not predict any sort of asymmetry effect. Additionally, within the pre-test and post-test conditions, our model predicts that there will be ordering effects, with later stimuli being rated as more similar than earlier stimuli; these effects, however, should average out given randomization of stimuli across participants.

As we have mentioned before, the different strengths of the subsymbolic links between two colors can cause asymmetries to arise in the degree to which they prime each other. Consider, on an intuitive level, Figure 1(b). Now, imagine that a model with this subsymbolic structure has the goal to judge how similar *blue5* and *blue1* are. First, both color chunks receive equal source activation from the goal chunk due to their equivalent co-occurrence links with it. They then provide each other with spreading activation according to the appropriate link strength. Here, as its greater number of links implies, *blue1* has been needed more times than *blue5*; this means that the link *blue1*→*blue5* is weaker than the link *blue5*→*blue1*. Therefore, *blue1* will receive more spreading activation than *blue5*, leading to an asymmetry in their similarity rankings.

As a result of this asymmetry, priming in the model predicts different effects for the pre- and post- tests. For the pre-test, the model predicts differences in similarity of forward and backward comparisons based solely on ordering effects of the stimuli. Given enough participants, these ordering effects average out over time to result in equal pre-test forward and backward comparisons. For the post-test, the model predicts that less frequently shown colors will spread more activation to more frequently shown colors than vice versa; i.e., it predicts that colors in forward comparisons will be ranked as more similar than those in backward comparisons.

Finally, our model predicts that the green ratings will be higher overall than the blue ratings, as well as that the 4-5 group’s ratings will be higher overall than the 1-2 group’s ratings. This is because of the specific hues chosen and is a purely perceptual point. The color similarity values do not differ depending on the direction of the comparison, or on whether the test is a pre- or post-test.

## Model Fit

In addition to the experimental results published in the original article (Polk et al., 2002), we also examined more detailed aggregate data provided to us by the authors. The data included the averages, for each subject, of ratings for trials of each condition (e.g., the average rating for each subject of all pre-test forward trials of blue hue, etc.). Since our model

is sensitive to the order in which stimuli are presented, we would have preferred to replicate the experiment exactly, including presenting the stimuli in the same order as in the original experiment. Because this information was not available, we instead used our model to simulate data from 1000 participants, in order to allow effects to better converge on the model's true predictions.

Our measurement of the model fit is done in two steps. First, the model needs to transform the similarity measures into an overall similarity rating. We do this post-hoc by fitting a linear regression model to the data, with the perceptual similarity value and total activation as the explanatory variables and the human participants' ratings as the dependent variable. We use total activation to maintain cognitive fidelity; it is unclear whether human minds can separately consider base-level and spreading activation values during cognitive tasks. Individual data points were the different conditions (e.g., the average rating across all participants in the 1-2 group of pre-test forward trials of blue hue, etc.). The model only considers the main effects of the two variables; this is because our goal is to show that the two similarity measures are the primary components of similarity ratings in this task, not to make any claims about how they are combined by the brain into a numerical rating. We take this approach because there are very few theories or accepted practices of how to convert continuous, numerical data to a rating scale.

Second, with this step completed, we compared the model's predicted ratings with the human participants' ratings. The model does indeed produce the two main significant effects of the human subject data, showing both a directional asymmetry in post-test comparisons as well as an increase in similarity ratings overall in the post-test. Figure 2 shows graphs of the numerical results for both the model and the human subject data. Note that error bars for the human data are not available.  $R^2$  (multiple) for these graphs across all data was 0.99; for blue only, 0.96; for green only, 0.96; for the 1-2 group (which saw *blue1*, *blue2*, *green1* and *green2* more frequently), 0.98, and for the 4-5 group (which saw *blue4*, *blue5*, *green4* and *green5* more frequently), 0.91.

For differences in color, the model's results, where blue pre-test ratings are slightly lower than green pre-test ratings, do not match well with the data's trends. The model does have overall higher ratings for the 4-5 group pre-test than the 1-2 group pre-test, but not to the extent of the data. The effects do not present a difficulty for the model, however, in large part because the experiment did not find these effects to be significant, presumably due to its small sample size of 35.

In terms of their individual contribution to explaining the data, color alone yields an  $R^2$  of 0.09; color and base-level activation produce an  $R^2$  of 0.79; and color, base-level activation and spreading activation an  $R^2$  of 0.99. This is intuitive and consistent with our account of the data: base-level activation is responsible for the larger main effect between the pre- and post-tests, while spreading activation correlates to the more modest interaction effect of forward versus back-

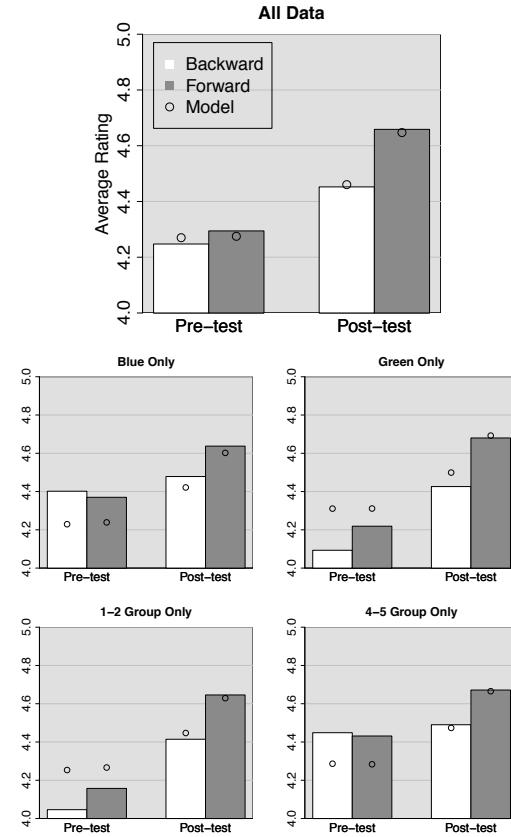


Figure 2: Main Results.

ward comparisons.

## Discussion

In this paper, we have introduced an account for similarity ratings that combines familiarity, priming, and perceptual similarity into a single judgment. We match our account to human subject data involving simple perceptual stimuli from Polk et al. (2002). Our approach explains that ratings rise over time because participants become more familiar with the stimuli, in general. Priming explains the asymmetry effect found because, inherently, low frequency concepts prime high frequency concepts more than the opposite. Finally, although the experiment did not reveal any significant differences between colors, our approach predicts that different pairs of colors will be slightly more similar than others due to pure perceptual similarity. Using these mechanisms, we show an excellent match to the human data.

Our approach is significant for at least two main reasons. First, we provide explicit, process-level mechanisms for determining similarity that explain, in a sense, how others' work (Rosch, 1975; Tversky, 1977) is realized by the human mind. Second, and perhaps more importantly, the mechanisms we offer as the basis of similarity have been shown to be pre-existing mechanisms in cognition that are also used for other

cognitive processes such as retrieval of memories, categorization, and problem solving (Anderson, 2007; Altmann & Trafton, 2002). This strengthens our approach since it also explains the pervasiveness of similarity in human cognition that has been found by a plethora of other research.

Finally, it is worth noting that familiarity, priming and perceptual similarity are not intended to be characterized as the ultimate, and only, way to determine similarity ratings. While they work well with simple perceptual stimuli, and we expect that their success will also extend to more complicated perceptual stimuli and simple concepts, we recognize that more complicated mechanisms are likely at work in, for example, the similarity of complex perceptual scenes, or the similarity of two short stories. We believe that such judgments likely involve some sort of structure alignment process as others have hypothesized for similarity judgments of higher-level stimuli or concepts (Markman, 1999; Goldstone, 1994). Instead, this paper is intended to introduce familiarity, priming and perceptual similarity as the foundation for similarity which other mechanisms can augment.

### Acknowledgments

This work was supported by the Office of Naval Research, grant numbers N0001412WX30002 and N0001411WX20516 to GT. The views and conclusions contained in this document do not represent the official policies of the U.S. Navy.

Many thanks to Thad Polk for making his data available to us. We would also like to acknowledge Wallace E. Lawson and Anthony M. Harrison for their helpful discussions during this work.

### References

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, 26, 39-83.

Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 22(3), 261-295.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford University Press.

Anderson, J. R., Albert, M. V., & Fincham, J. M. (2005, August). Tracing problem solving in real time. *Journal of Cognitive Neuroscience*, 17, 1261-1274.

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, 38(4), 341-380.

Breslow, L. A., Ratwani, R. M., & Trafton, J. G. (2009). Cognitive models of the influence of color scale on data visualization tasks. *Human Factors*, 51(3), 321-338.

Breslow, L. A., Trafton, J. G., & Ratwani, R. M. (2009). A perceptual process approach to selecting color scales for complex visualizations. *Journal of Experimental Psychology: Applied*, 15(1), 25-34.

CIE. (2001). *Improvements to industrial colour-difference evaluation* (Tech. Rep. No. CIE 142-2001). Vienna: Commission Internationale de l'Eclairage.

Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, 97, 3-18.

Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 20(3-28).

Harrison, A. M., & Trafton, J. G. (2010). Cognition for action: an architectural account for "grounded interaction". In *Proceedings of the Annual Conference on Cognitive Science*.

Kozima, H., & Furugori, T. (1993). Similarity between words computed by spreading activation on an English dictionary. In *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics*. Stroudsburg, PA, USA.

Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100(2), 254-278.

Medin, D. L., Goldstone, R. L., & Markman, A. B. (1995). Comparison and choice: Relations between similarity processes and decision processes. *Psychonomic Bulletin and Review*, 2, 1-19.

Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (p. 363-393). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.

Novick, L. R. (1990). Representational transfer in problem solving. *Psychological Science*, 128-132.

Polk, T. A., Behensky, C., Gonzalez, R., & Smith, E. E. (2002). Rating the similarity of simple perceptual stimuli: Asymmetries induced by manipulating exposure frequency. *Cognition*, 82(3), B75-B88.

Roediger, H. L. (1990). Implicit memory: Retention without remembering. *American Psychologist*, 45, 1043-1056.

Rorissa, A. (2005). *Perceived features and similarity of images: An investigation into their relationships and a test of tversky's contrast model*. Unpublished doctoral dissertation, University of North Texas.

Rosch, E. (1975). Cognitive reference points. *Cognitive Psychology*, 7, 532-547.

Schneider, D. W., & Anderson, J. R. (2011). A memory-based model of hick's law. *Cognitive Psychology*, 62(3), 193-222.

Smith, L. B., & Heise, D. (1992). Perceptual similarity and conceptual structure. *Advances in Psychology – Percepts, Concepts and Categories: The Representation and Processing of Information*, 93, 233-272.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.

Ulhaque, E., & Bahn, K. D. (1992). A spreading activation model of consumers' asymmetric similarity judgment. *Advances in Consumer Research*, 19, 782-786.